REMARKS

The non-final Office Action, dated April 20, 2006, considered and rejected claims 1-21. Claims 1, 6-8, 13-15, 20 and 21 were rejected under 35 U.S.C. § 102(b) as being unpatentable over Brown (U.S. Patent No. 6,101,585). Claims 2, 5, 9, 12, 16 and 19 were rejected under 35 U.S.C. § 103(a) as being unpatentable over Brown in view of Mani-Meitav (U.S. Publ. No. 2005/0216788). Claims 3, 4, 10, 11 17 and 18 were rejected under 35 U.S.C. § 103(a) as being unpatentable over Brown in view of Mani-Meitav, and further in view of Lee (U.S. Patent No. 6,078,930).[1]

By this paper, claims 1, 2, 8 and 15 have been amended, and claim 22 has been added, while no claims have been cancelled.[2] Accordingly, following this paper, claims 1-22 are pending, of which claims 1, 8 and 15 are the only independent claims at issue.

As reflected in the claims listing above, the present invention is generally directed to computer-readable media, methods and systems for backing-up and restoring a data set of a node in a distributed system. As recited in claim 1, for example, a computer-readable medium includes computer-executable instructions for performing various acts. For instance, the computer-executable instructions initiate a back-up operation and set state data of the data set to indicate that the data set has been backed up. Further, a snapshot of the contents of the data set is taken to generate a back-up copy. The contents of the data set of which the snapshot is taken include the previously set state data indicating that the data set is backed-up, and a plurality of changes that are each identified by a replication identification number of the node and a serial number. The state data of the data set is then reset after taking the snapshot to indicate that the data set is in normal operation.

Claims 8 and 15 recite a method and system, respectively which include generally correspond to the computer-readable medium recited in claim 1. In addition, claim 15 further recites a replication server for replicating changes made to the data set to other nodes in a distributed system.

---

[1] Although the prior art status of the cited art is not being challenged at this time, Applicant reserves the right to challenge the prior art status of the cited art at any appropriate time, should it arise. Accordingly, any arguments and amendments made herein should not be construed as Applicant acquiescing to any prior art status of the cited art.

[2] Support for the amendments and new claims can be found throughout Applicant's disclosure, including, by way of representation only, the discussion in paragraphs 7 and 22-24, as well as in Figure 4, of the originally filed application.

While the Brown reference is generally directed towards a system for backing-up files of a computer system, Applicant respectfully submits that the cited art fails to anticipate or make obvious the claimed invention. For example, the cited art fails to disclose or suggest, among other things, a computer-readable medium (claim 1), method (claim 8) or system (claim 15) in which state data *of the data set* is set to indicate that the data set is in a backed-up state and taking a snapshot of the contents of the data set, *including the set state data* indicating the data set is in a backed-up state, to generate a backup copy, and thereafter resetting the state data of the data set to indicate that the data set is in normal operation, as claimed in combination with the other recited claim elements. In fact, in direct contrast to the present invention, in which the state data of the data set is modified and included in the snapshot when taken, Brown appears to disclose a system in which state data of the data set is unmodified when the snapshot is taken, and wherein the state data *within the snapshot* is modified, rather than the contents of the data set.

In particular, Brown discloses a backup mechanism for incremental backup operations that properly backup files that are modified concurrently with the backup operation. (Col. 3, ll. 15-23). As disclosed in Brown, files are stored in a file system (or container) which has its own ID number. (Col. 4, ll. 65-67). In addition, each file is associated with an archive bit change number ("ABCN") attribute. (Col. 6, ll. 5-10). The ABCN attribute is used to identify changes to files, and is incremented each time the associated file is modified. (Col. 6, ll. 15-18).

When a back-up operation is performed, the contents of a data container are duplicated as a read-only, snapshot container which is virtually an identical copy of the original container. (Col. 5, ll. 6-17). During the creation of the snapshot, however, the state of most file attributes remains the same. (Col. 5, ll. 48-51). The container ID associated with each file, however, does not remain the same. In particular, the file system converts the container ID *in the snapshot* to reflect that the files are on the snapshot rather than the original container. (Col. 5, ll. 52-54; Col. 6, ln. 5). Subsequently, when the file system wishes to indicate that a backup has been made, it converts the container ID *on the shapshot* so that it can compare ABCN numbers and determine whether a file has been updated since the backup was initiated. (Col. 6, ll. 26-40).

Accordingly, Brown teaches updating ABCN file attributes of files in a file system/container when the file is modified, and updating the container ID of the snapshot during a backup operation. Brown fails, however, to disclose wherein the *state data of the data set* itself

is set to indicate that the data set is backed up and in which the snapshot of the contents of the data set includes the state data set of the data set to indicate that the data set is backed up. In fact, Brown fails to disclose, for example, that the Container ID or some type of state data of *the original data set* is set or reset to indicate that the data set is backed up, and particularly fails to disclose that it is done such that the set state data indicating the data is backed-up is carried over into the snapshot and thereafter reset, as claimed. In fact, it appears that the Container ID *of the original data set* remains the same before, during, and after the backup operation in Brown, such that it does not distinguish whether the data is or is not backed up. In fact, only the Container ID *in the snapshot* appears to ever be changed.

As a result, Brown fails to teach that state data *of the data set* is: (i) set to indicate a data set is backed up; and (ii) carried into a snapshot in a form that indicates the data set is backed up, as claimed in combination with the other recited claim elements. Moreover, inasmuch as the state data of the original data set in Brown is unchanged, Brown necessarily also fails to disclose resetting the state data *of the data set* to indicate that the data set is in normal operation after taking the snapshot, as claimed in combination with the other recited claim elements.

Accordingly, and for at least these reasons, Applicant respectfully that the independent claims are allowable over the cited references. Moreover, inasmuch as each independent claim is allowable over the cited references, it will be appreciated that all other rejections of record with respect to the dependent claims are now moot and need not be addressed individually. Nevertheless, to further differentiate between the cited references and the present invention, various dependent claims will be specifically addressed.

With respect to claims 2, 9 and 16, Applicant respectfully submits that the cited art fails to disclose or suggest embodiments that include changing the replication identification number of the node from an old value used before the backup operation to a new value in response to detecting that the restored data set indicates that the data set is in the backed up state, as recited in combination with the other claim elements. For this teaching, the Office Action states that it would be inherent that when a node is brought back online a new IP/GUID would be assigned. Applicant respectfully disagrees.

Applicant notes that there are various ways a computer system may operate with regard to its identification number such that it is not necessary that the IP/GUID be newly assigned. For example, Applicant notes that it the node may operate with a static IP/GUID such that it

maintains the same IP/GUID. Moreover, Applicant notes that Brown does not have any disclosure even directed to the use of an IP/GUID.

In addition, as recited in the pending claims, the replication identification number is used to identify each of a plurality of changes in a data set. Brown does not disclose that the IP/GUID is used to identify changes. In fact, the Office Action states that the file system ID reads on the claimed "replication identification number." Significantly, however, Brown does not teach that the file system ID of a node is ever changed. In fact, the snapshot file system ID is changed to match the original file system ID such that restoration using a backup would necessarily result in the same file system ID as in the original.

With respect to claims 3, 10 and 17, Applicant respectfully submits that the cited art fails to teach embodiments for storing a lowest uncommitted serial number corresponding to a lowest serial number of changes made to the data set that are not yet committed prior to taking the snapshot, as recited in combination with the other recited claim elements. The Office Action acknowledges that Brown and Mani-Meitav fail to disclose such a teaching. To remedy this failure, the Office Action relies on the disclosure in the Lee reference related to the use of a recovery timestamp value. (Col. 4, ll. 29-35). Applicant notes, however that the recovery timestamp value is an indication, from a logical clock, of the relative time at which changes were last performed. (Col. 4, ll. 24-28). Accordingly, Lee discloses use of a timestamp to indicate when changes were last made. The timestamp does not correspond to a lowest one of a serial number of changes made that are not yet committed. In fact, Lee does not appear to have any disclosure regarding committed or uncommitted changes.

With respect to new claim 22, Applicant also respectfully submits that the cited art fails to disclose or suggest a system, method or computer-readable medium in which state data that is set and reset is a collective value for the data set. In fact, and to the contrary, Brown teaches that each file that is backed-up includes a separate file system ID (or container ID) which is then changed in the snapshot of the file system or container. In other words, Brown teaches that the set of data being backed-up includes multiple file system IDs that operate on individual files, rather than collectively for an entire data set, as claimed in combination with the other recited claim elements.

In view of the foregoing, Applicant respectfully submits that the other rejections to the claims are now moot and do not, therefore, need to be addressed individually at this time. It will

be appreciated, however, that this should not be construed as Applicant acquiescing to any of the purported teachings or assertions made in the last action regarding the cited art or the pending application, including any official notice. Instead, Applicant reserves the right to challenge any of the purported teachings or assertions made in the last action at any appropriate time in the future, should it arise. Furthermore, to the extent that the Examiner has relied on any Official Notice, explicitly or implicitly, Applicant specifically requests that the Examiner provide references supporting the teachings officially noticed, as well as the required motivation or suggestion to combine references with the other art of record.

In the event that the Examiner finds remaining impediment to a prompt allowance of this application that may be clarified through a telephone interview, the Examiner is requested to contact the undersigned attorney.

Dated this 20th day of July, 2006.

Respectfully submitted,

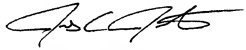

RICK D. NYDEGGER
Registration No. 28,651
JENS C. JENKINS
Registration No. 44,803
COLBY C. NUTTALL
Registration No. 58,146
Attorneys for Applicant
Customer No. 047973

JCJ:CCN:ahy